

UNITED STATES PATENT APPLICATION

FOR

METHOD AND APPARATUS FOR REPRESENTING LABEL SWITCHED PATHS

Inventors:

**Rahul Aggarwal
Ravi Chandra**

Prepared by:

Blakely, Sokoloff, Taylor & Zafman LLP
12400 Wilshire Blvd., Suite 700
Los Angeles, California 90025
(408) 720-3800

EXPRESS MAIL CERTIFICATE OF MAILING

"Express Mail" mailing label number: EL 863955615 US Date of Deposit: 12-31-2001

I hereby certify that I am causing this paper or fee to be deposited with the United States Postal Service
"Express Mail Post Office to Addressee" service on the date indicated above and that this paper or fee has been
addressed to the Commissioner of Patents and Trademarks, Washington, D. C. 20241

SHENISE RAMDEEN

(Typed or printed name of person mailing paper or fee)

(Shenise Ramdeen)

(Signature of person mailing paper or fee)

12-31-2001

(Date signed)

METHOD AND APPARATUS FOR REPRESENTING LABEL SWITCHED PATHS

BACKGROUND OF THE INVENTION

Field of the Invention

[0001] The invention relates to the field of communication. More specifically, the invention relates to communication networks.

Background of the Invention

[0002] The multi-protocol label switching (MPLS) protocol may be categorized as a network layer protocol of the Open Standards Institute (OSI) reference model. MPLS provides a method for generically tunneling data through networks with label switched paths (LSPs).

[0003] Routers typically internally represent each connection (whether it be an IP route, a label switched path, etc.) as an interface or set of interfaces, which is a network layer entity. Since an interface is a network layer entity, it includes various pieces of information needed for the network layer.

[0004] Figure 1 (Prior Art) is a diagram illustrating an exemplary data structure for an interface. An interface structure 101 includes multiple fields describing the interface. An interface ID field 103 indicates a value identifying the interface. An interface type field 105 describes the type of interface (e.g., Ethernet, ATM, PoS, etc.). An IP address field 107 identifies a 32-bit IP address corresponding to the interface. A secondary IP address field 109 indicates a second 32-bit IP address for the interface. A maximum transmission unit (MTU) field 111 indicates the maximum allowable packet size to be transmitted with the interface. A bandwidth field 113 indicates the amount of bandwidth allocated to the interface. The interface structure 101 is a relatively expensive structure.

[0005] The relatively expensive interface structure consumes memory of a network element and consumes bus resources when the network element transfers interface structures to its line cards. A network element cannot maintain a relatively large number of label switched paths without hampering its performance because of the relatively expensive cost of the interface structure.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65
66
67
68
69
70
71
72
73
74
75
76
77
78
79
80
81
82
83
84
85
86
87
88
89
90
91
92
93
94
95
96
97
98
99
100
101
102
103
104
105
106
107
108
109
110
111
112
113
114
115
116
117
118
119
120
121
122
123
124
125
126
127
128
129
130
131
132
133
134
135
136
137
138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161
162
163
164
165
166
167
168
169
170
171
172
173
174
175
176
177
178
179
180
181
182
183
184
185
186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
201
202
203
204
205
206
207
208
209
210
211
212
213
214
215
216
217
218
219
220
221
222
223
224
225
226
227
228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259
260
261
262
263
264
265
266
267
268
269
270
271
272
273
274
275
276
277
278
279
280
281
282
283
284
285
286
287
288
289
290
291
292
293
294
295
296
297
298
299
300
301
302
303
304
305
306
307
308
309
310
311
312
313
314
315
316
317
318
319
320
321
322
323
324
325
326
327
328
329
330
331
332
333
334
335
336
337
338
339
340
341
342
343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377
378
379
380
381
382
383
384
385
386
387
388
389
390
391
392
393
394
395
396
397
398
399
400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415
416
417
418
419
420
421
422
423
424
425
426
427
428
429
430
431
432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485
486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539
540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593
594
595
596
597
598
599
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647
648
649
650
651
652
653
654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701
702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809
810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863
864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917
918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971
972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025
1026
1027
1028
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1040
1041
1042
1043
1044
1045
1046
1047
1048
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079
1080
1081
1082
1083
1084
1085
1086
1087
1088
1089
1090
1091
1092
1093
1094
1095
1096
1097
1098
1099
1100
1101
1102
1103
1104
1105
1106
1107
1108
1109
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1130
1131
1132
1133
1134
1135
1136
1137
1138
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187
1188
1189
1190
1191
1192
1193
1194
1195
1196
1197
1198
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241
1242
1243
1244
1245
1246
1247
1248
1249
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295
1296
1297
1298
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349
1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
1372
1373
1374
1375
1376
1377
1378
1379
1380
1381
1382
1383
1384
1385
1386
1387
1388
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403
1404
1405
1406
1407
1408
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457
1458
1459
1460
1461
1462
1463
1464
1465
1466
1467
1468
1469
1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1490
1491
1492
1493
1494
1495
1496
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511
1512
1513
1514
1515
1516
1517
1518
1519
1520
1521
1522
1523
1524
1525
1526
1527
1528
1529
1530
1531
1532
1533
1534
1535
1536
1537
1538
1539
1540
1541
1542
1543
1544
1545
1546
1547
1548
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565
1566
1567
1568
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1579
1580
1581
1582
1583
1584
1585
1586
1587
1588
1589
1590
1591
1592
1593
1594
1595
1596
1597
1598
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619
1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1630
1631
1632
1633
1634
1635
1636
1637
1638
1639
1640
1641
1642
1643
1644
1645
1646
1647
1648
1649
1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1670
1671
1672
1673
1674
1675
1676
1677
1678
1679
1680
1681
1682
1683
1684
1685
1686
1687
1688
1689
1690
1691
1692
1693
1694
1695
1696
1697
1698
1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1720
1721
1722
1723
1724
1725
1726
1727
1728
1729
1730
1731
1732
1733
1734
1735
1736
1737
1738
1739
1740
1741
1742
1743
1744
1745
1746
1747
1748
1749
1750
1751
1752
1753
1754
1755
1756
1757
1758
1759
1760
1761
1762
1763
1764
1765
1766
1767
1768
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1780
1781
1782
1783
1784
1785
1786
1787
1788
1789
1790
1791
1792
1793
1794
1795
1796
1797
1798
1799
1800
1801
1802
1803
1804
1805
1806
1807
1808
1809
1810
1811
1812
1813
1814
1815
1816
1817
1818
1819
1820
1821
1822
1823
1824
1825
1826
1827
1828
1829
1830
1831
1832
1833
1834
1835
1836
1837
1838
1839
1840
1841
1842
1843
1844
1845
1846
1847
1848
1849
1850
1851
1852
1853
1854
1855
1856
1857
1858
1859
1860
1861
1862
1863
1864
1865
1866
1867
1868
1869
1870
1871
1872
1873
1874
1875
1876
1877
1878
1879
1880
1881
1882
1883
1884
1885
1886
1887
1888
1889
1890
1891
1892
1893
1894
1895
1896
1897
1898
1899
1900
1901
1902
1903
1904
1905
1906
1907
1908
1909
1910
1911
1912
1913
1914
1915
1916
1917
1918
1919
1920
1921
1922
1923
1924
1925
1926
1927
1928
1929
1930
1931
1932
1933
1934
1935
1936
1937
1938
1939
1940
1941
1942
1943
1944
1945
1946
1947
1948
1949
1950
1951
1952
1953
1954
1955
1956
1957
1958
1959
1960
1961
1962
1963
1964
1965
1966
1967
1968
1969
1970
1971
1972
1973
1974
1975
1976
1977
1978
1979
1980
1981
1982
1983
1984
1985
1986
1987
1988
1989
1990
1991
1992
1993
1994
1995
1996
1997
1998
1999
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2026
2027
2028
2029
2030
2031
2032
2033
2034
2035
2036
2037
2038
2039
2040
2041
2042
2043
2044
2045
2046
2047
2048
2049
2050
2051
2052
2053
2054
2055
2056
2057
2058
2059
2060
2061
2062
2063
2064
2065
2066
2067
2068
2069
2070
2071
2072
2073
2074
2075
2076
2077
2078
2079
2080
2081
2082
2083
2084
2085
2086
2087
2088
2089
2090
2091
2092
2093
2094
2095
2096
2097
2098
2099
2100
2101
2102
2103
2104
2105
2106
2107
2108
2109
2110
2111
2112
2113
2114
2115
2116
2117
2118
2119
2120
2121
2122
2123
2124
2125
2126
2127
2128
2129
2130
2131
2132
2133
2134
2135
2136
2137
2138
2139
2140
2141
2142
2143
2144
2145
2146
2147
2148
2149
2150
2151
2152
2153
2154
2155
2156
2157
2158
2159
2160
2161
2162
2163
2164
2165
2166
2167
2168
2169
2170
2171
2172
2173
2174
2175
2176
2177
2178
2179
2180
2181
2182
2183
2184
2185
2186
2187
2188
2189
2190
2191
2192
2193
2194
2195
2196
2197
2198
2199
2200
2201
2202
2203
2204
2205
2206
2207
2208
2209
2210
2211
2212
2213
2214
2215
2216
2217
2218
2219
2220
2221
2222
2223
2224
2225
2226
2227
2228
2

BRIEF SUMMARY OF THE INVENTION

[0006] A method and apparatus for representing label switched paths is described.

According to one aspect of the invention, a computer implemented method provides for representing a label switched path (LSP) with a data structure. The data structure indicates a virtual ingress and an identifier to distinguish the LSP from other LSPs represented with the same virtual ingress.

[0007] These and other aspects of the present invention will be better described with reference to the Detailed Description and the accompanying Figures.

BRIEF DESCRIPTION OF THE DRAWINGS

[0008] The invention may best be understood by referring to the following description and accompanying drawings that are used to illustrate embodiments of the invention. In the drawings:

[0009] Figure 1 (Prior Art) is a diagram illustrating an exemplary data structure for an interface.

[0010] Figure 2 is an exemplary diagram of a forwarding data structure according to one embodiment of the invention.

[0011] Figure 3 is an exemplary diagram illustrating interaction between a control plane and a data plane according to one embodiment of the invention.

[0012] Figure 4 is an exemplary diagram of the LSP structure 208A according to one embodiment of the invention.

[0013] Figure 5A is an exemplary diagram illustrating the LFIB 215A according to one embodiment of the invention.

[0014] Figure 5B is an exemplary diagram illustrating the LSP adjacency table 216A according to one embodiment of the invention.

[0015] Figure 6 is an exemplary diagram illustrating the data plane 221 according to one embodiment of the invention.

[0016] Figure 7A is an exemplary diagram illustrating a packet transmitted along an LSP at a network element acting as an ingress router according to one embodiment of the invention.

[0017] Figure 7B is an exemplary diagram illustrating transmission of a packet along an LSP at a network element acting as an egress router according to one embodiment of the invention.

[0018] Figure 7C is an exemplary diagram illustrating transmission of a packet along an LSP at a network element acting as a transient router according to one embodiment of the invention.

[0019] Figure 8 is a diagram of an exemplary network element according to one embodiment of the invention.

DETAILED DESCRIPTION OF THE INVENTION

[0020] In the following description, numerous specific details are set forth to provide a thorough understanding of the invention. However, it is understood that the invention may be practiced without these specific details. In other instances, well-known circuits, structures, standards, and techniques have not been shown in detail in order not to obscure the invention.

[0021] A method and apparatus for representing label switched paths is described.

According to one embodiment of the invention, a forwarding data structure is used. These forwarding data structures are used for processing LSP packets and non-LSP packets. These forwarding data structures are less expensive than interface structures because they do not contain network layer information. The forwarding data structures contain information to describe links from a network element to its neighboring network elements.

[0022] Figure 2 is an exemplary diagram of a forwarding data structure according to one embodiment of the invention. In Figure 2, forwarding structure 201 includes the following fields: a slot field 203, a port field 205, and a flow identifier field 207. The forwarding data structure may also include a forwarding feature field 209. The slot field 203 indicates a slot of a network element. The port field 205 indicates a port of the slot indicated in the slot field 203. However, a port may not always be identified (e.g., because a card in a slot only has a single egress port). The flow identifier field 207 indicates an identifier for distinguishing different flows of traffic received on the slot and port indicated in the forwarding data structure 201. The forwarding features field 209 indicates forwarding features (e.g., packet counters, quality of service, packet classification) to be applied to packets. In alternative embodiments, forwarding features are located in a separate data structure and indexed by a slot, port, and flow identifier. A number of fields in a number of structures and tables will be described herein. It should be understood that each such field can be implemented in

different embodiments to “indicate” information using different techniques (e.g., store the information directly in the field, store a reference to the information, etc.).

[0023] In a network element with the forwarding structure, non-LSP packets are processed with a forwarding structure for ingress into the network element and a forwarding structure for egress from the network element. Forwarding features may be applied to non-LSP packets with the forwarding structure 201 and switched through a network element with the forwarding structure. Similarly, forwarding features may be applied to LSP packets associated with a forwarding structure that described an LSP. Various embodiments may indicate LSPs with forwarding structures differently. For example, all forwarding structures that describe LSPs may indicate the same virtual slot and virtual port, but different flow identifiers to distinguish LSPs. Alternatively, forwarding structures that describe LSPs may indicate the same virtual slot for all LSPs, but distinguish LSPs with different virtual ports.

[0024] The described forwarding structure is less expensive than an interface structure and may be used to represent LSPs. Representing LSPs with such a relatively inexpensive data structure provides substantial memory savings and enables the support of relatively large numbers of LSPs without reducing performance of a network element. In addition, representing LSPs with such relatively inexpensive structures reduces the amount of resources consumed when downloading LSP information to the control cards of a network element.

[0025] The control plane of a network element may be implemented differently to represent LSPs with forwarding structures. In one embodiment, functionality for representing LSPs with forwarding structures is added to the module that manages interface structures. In another embodiment, a separate module is implemented with the functionality to support forwarding structures for LSPs and non-LSP packets. In an alternative embodiment, a module is implemented with the functionality to manage forwarding

structures for non-LSP packets and a separate module is implemented to manage forwarding structures for LSPs. Another embodiment will be illustrated in Figure 3.

[0026] Figure 3 is a diagram illustrating an exemplary control plane according to one embodiment of the invention. In Figure 3, a control plane 301 includes the following: signaling protocol modules 303A-303C, a static LSP module 305, a label manager 307, an interface manager 309, routing protocol modules 211A-211J, and a routing information base 313. The signaling protocol modules 303A-303C (e.g., label distribution protocol, RSVP, BGP label piggybacking, etc.) establish label switched paths (LSPs) and establish labels with the label manager 307. In addition to the signaling protocol modules 303A-303C, the static LSP module 305 establishes LSPs configured by a user. The user may configure the LSPs via a user interface, a script, etc. The static LSP module 305 also establishes labels with the label manager 307 for configured LSPs. The label manager 307 manages LSP structures 308A-308W.

[0027] Figure 4 is an exemplary diagram of the LSP structure 308A according to one embodiment of the invention. In Figure 4, the LSP structure 308A includes multiple fields describing an LSP. A forwarding structure field 401 indicates a forwarding structure for an LSP. The LSP structure 308A also includes an egress label field 403. The egress label field 403 indicates an egress label identifier to be applied to a packet transmitted out of a network element. An LSP ID field 405 indicates a value identifying the LSP. Various embodiments may allocate the LSP ID differently. In one embodiment, signaling protocol modules 303A-303C and the static LSP module 305, generate the LSP ID. In another embodiment, the label manager 307 generates the LSP ID. The LSP structure 308A also includes an adjacency field 407. The adjacency field 407 indicates adjacency information. In one embodiment, the adjacency information includes an adjacency ID, a forwarding structure for an egress from a network element, and encapsulation information, which will be described later in more detail.

In another embodiment, the adjacency information includes an encapsulation ID to reference the encapsulation information. In an alternative embodiment, adjacency information includes a slot, an adjacency ID, and encapsulation information.

[0028] In Figure 4, the LSP structure 308A also includes a port field 409. The port field 409 is used to identify a port to transmit packets along the LSP identified by the LSP structure 308A. However, a port may not be identified because a card in a slot only has a single egress port. In certain embodiments, the port field 409 may not be included in the LSP structure 408A because the adjacency field 407 indicates adjacency information that includes the port.

[0029] Returning to Figure 3, the label manager 307 published forwarding structures to the interface manager 309. In one embodiment, the label manager 307 manages all forwarding structures for a network element. In another embodiment, the label manager manages forwarding structure allocated for LSPs while the interface manager manages forwarding structure for non-LSP traffic. The interface manager 309 may publish forwarding structures for LSPs to the routing protocol modules 311A – 311J. In certain situations, a routing protocol module will use an LSP for routing purposes, but in more situations, the routing protocol modules 311A – 311J do not use LSPs. Hence, resources are used inefficiently if the interface manager published all forwarding structures for LSPs to the routing protocol modules 311A – 311J.

[0030] The label manager 307 publishes forwarding structures for LSPs to the interface manager 309, but the interface manager 309 only publishes forwarding structures for LSPs to routing protocol modules that have registered interest in forwarding structures for LSPs. In an alternative embodiment, the label manager 307 publishes forwarding structures for LSPs to the interface manager 309 with indicators. The indicator for each forwarding structure either permits or inhibits the interface manager's 309 ability to publish each forwarding

structure. In an alternative embodiment, the routing protocol modules 311A – 311J register interest in certain LSPs with the label manager 307. The label manager 307 may publish those certain LSPs to the registered ones of the routing protocol modules 311A – 311J or publish those certain LSPs to the interface manager 309, which then publishes all forwarding structures received from the label manager 307.

[0031] The routing protocol modules 311A-311J publish information from their routing tables to the routing information base 313. The information includes the interfaces and/or forwarding structures in the routing tables of the routing protocol modules 311A-311J. The routing information base 313 publishes information as forwarding information bases (FIBs) 317A-317D to a data plane 321.

[0032] The label manager 307 distributes information from the LSP structures 308A-308W to the data plane 321. The label manager 307 distributes information from the LSP structures 308A-308W to label forwarding information bases (LFIBs) 315A-315C and LSP adjacency tables 316A-316D. The label manager 307 may target distribution of information to certain LFIBs 315A-315C and certain LSP adjacency tables 316A-316D.

[0033] When establishing an LSP, the label manager 309 determines either an ingress (i.e., slot, port, and flow identifier) for the LSP into the network element, an egress for the LSP from the network element, or both, depending on the role of the network element (i.e., whether the network element is acting as an egress router, an ingress router, or a transient router). With this information, the label manager 309 distributes information to one of the LFIBs 315A – 315C located on the line card determined as the ingress point into the network element for the LSP and distributes information to the one of the LSP adjacency tables 316A – 316D located on the line card determined to be the egress points from the network element for the LSP

[0034] Representing an LSP with the forwarding structure instead of an interface reduces memory consumption and reduces the amount of data communicated across from the control plane 301 to the data plane 321. The LSP structure is a relatively inexpensive structure and identifies information for transmission of packets corresponding to the LSP identified by the LSP structure. The reduced cost of representing an LSP enables a network element to maintain a relatively large number of LSPs. Resources are further conserved because forwarding structures are not propagated to all routing protocol modules, the routing information base, and all forwarding information bases. In addition, representing an LSP with a forwarding structure provides uniform presentation of forwarding structures for LSPs and forwarding structures for non-LSP traffic to a network element.

[0035] Figure 5A is an exemplary diagram illustrating the LFIB 315A according to one embodiment of the invention. In Figure 5A, the LFIB 315A includes two fields for each entry: a forwarding structure field 501 and an adjacency field 503. Although the LFIB 315A is illustrated as a table, the illustration is meant to be conceptual. The LFIB 315A may be implemented in a variety of ways (e.g., a binary search tree, a radix trie, a hash table, etc.). As previously described, the label manager 307 populates the LFIB 315A. When a packet is received that includes a label, the LFIB 315A is indexed into for an entry that corresponds to the label (the manner of this indexing depends on the implementation; while in one embodiment the LFIB is addressed by some form of the label, alternative embodiments may use other techniques (e.g., storing some form of the labels in each entry and searching the table using well known techniques, etc.)). The forwarding structure field 501 indicates a forwarding structure. The adjacency field 503 indicates adjacency information.

[0036] Figure 5B is an exemplary diagram illustrating the LSP adjacency table 316A according to one embodiment of the invention. Although the LSP adjacency table 316A is illustrated as a table, it may be implemented with a variety of data structures (e.g., a binary

search tree, a radix trie, a hash table, etc.). In Figure 5B, the LSP adjacency table 316A includes three fields for each entry: an encapsulation field 505, a port field 507, and a forwarding structure field 509. The encapsulation field 505 indicates an encapsulation that includes an egress label to be added to a set of packets. The port field 507 indicates an egress port. Alternative embodiments may not include an egress port if not necessary (e.g., a card with a single egress port). The forwarding structure field 509 indicates a forwarding structure for the LSP in order to apply forwarding features (e.g., packet counters and packet classification).

[0037] Figure 6 is an exemplary diagram illustrating the data plane 321 according to one embodiment of the invention. In Figure 6, the data plane 321 includes line cards 601A-601D. The line cards 601A-601C host LFIBS 315A-315C, FIBs 317A-317C, LSP adjacency tables 316A-316C, and adjacency tables 319A-319D. The line card 601D is illustrated as hosting the FIB 317D, the LSP adjacency table 316D, and the adjacency table 319D. The line card 601D does not host an LFIB in order to illustrate that a line card may not have an LFIB. Certain line cards may not have LFIBS and/or LSP adjacency tables because of the targeted distribution of information from the control plane. The label manager 307 populates the LFIBS 315A-315C on relevant line cards with relevant information from the LSP structures 308A-308W. As previously explained, the label manager 307 may determine relevant line cards with information received from the signaling protocols 303A-303C and/or the user interface 305. In contrast, the routing information base 313 downloads forwarding information to the FIBs 317A-317D on each of the line cards 601A-601D. Figure 6 also illustrates the label manager 307 populating the LSP adjacency tables 316A-316D respectively located on the line cards 601A-601D. Although the line card 601D does not host an LFIB, the label manager 307 may populate the LSP adjacency table 316D with information corresponding to an LSP if the network element is acting as an ingress router or

an egress router for an LSP as described later in more detail. The routing information base 313 populates the adjacency tables 319A-319D, but is not illustrated in order to avoid obscuring the invention.

[0038] As in the control plane, resources of the data plane are conserved because LSPs are represented with the forwarding structures. In addition, targeted distribution of LSP information from the control plane provides for further conservation of resources, to the point that a line card may not have any LSP information.

[0039] Figures 7A - 7C are exemplary diagrams illustrating transmission of a packet along an LSP according to one embodiment of the invention. Figure 7A is an exemplary diagram illustrating a packet transmitted along an LSP at a network element acting as an ingress router according to one embodiment of the invention. In Figure 7A, the line card 601D receives a packet 701A that has been decapsulated from a data link layer encapsulation and associated with a forwarding structure that indicates the ingress into the network element. The packet 701A includes a payload and an IP header. The payload may include additional encapsulation information. A forwarding engine 702D receives the packet 701A and processes the packet 701A in accordance with the FIB 317D. The forwarding engine 702D may also apply forwarding features (quality of service, packet counters, packet classification, etc.) to the packet 701A in accordance with the associated forwarding structure..

[0040] A switching medium 703 receives the packet 701A and adjacency information 709. The switching medium 703 forwards the packet 701A in accordance with the adjacency information 709. The switching medium 703 is hardware and/or software for determining where to forward packets. In one embodiment, the switching medium 703 is physically separate from the line cards. In another embodiment, the switching medium 703 is software and/or hardware on individual line cards. For example, the switching medium 703 may be software located on the line card 601D that processes the adjacency information 709 to

determine which line of a set of lines interconnecting all line cards to forward the packet 701A. The switching medium 703 then forwards the packet 701A and the adjacency information 709 to the line card 601A. In an alternative embodiment, the switching medium 703 may receive the packet 701A and an adjacency ID. The switching medium 703 then accesses adjacency information identified by the adjacency ID.

[0041] A forwarding engine 702A on the line card 701A receives the packet 701A and the adjacency information 709. The forwarding engine 702A processes the packet 701A in accordance with the LSP adjacency table 316A. The forwarding engine 702A uses the adjacency information 709 to look up encapsulation information and possibly a port for the packet 701B. The forwarding engine may also apply forwarding features associated with a forwarding structure that indicates the LSP. The forwarding engine 702A then transmits the packet 701B that includes an egress label identified in the LSP adjacency table 316A after encapsulating the packet 701B in a data link layer encapsulation.

[0042] Figure 7B is an exemplary diagram illustrating transmission of a packet along an LSP at a network element acting as an egress router according to one embodiment of the invention. In Figure 7B, the line card 601A receives a packet 701A that has been decapsulated from a data link layer encapsulation and that includes an ingress label and a payload. The payload may include additional encapsulations. The forwarding engine 702A receives the packet 701A and pops the label from the packet 701A. The forwarding engine 702A then processes the packet 701A in accordance with the FIB 317A and determines adjacency information from the FIB 317A. The forwarding engine 702A may apply forwarding features (e.g., quality of service, packet classification, packet counters, etc.), which are associated with the forwarding structure indicated in the FIB 317A, to the packet 701A. The forwarding engine 702A then passes the packet 701A and the identified adjacency information 709 to the switching medium 703.

[0043] The switching medium 703 forwards the packet 701A and the adjacency information 709 to the line card 601C, assuming the adjacency information identifies the slot hosting the line card 601C. A forwarding engine 702C on the line card 601C processes the packet 701A and the adjacency information 709. The forwarding engine 702C determines the encapsulation for the packet 701A in accordance with the encapsulation corresponding to the adjacency information 709. The forwarding engine 702C may apply forwarding features (e.g., packet classification, packet counters, etc.) associated with the forwarding structure identified by the adjacency information 709. The forwarding engine 702C may also determine the proper egress port for the packet 701A in accordance with the adjacency information 709. The forwarding engine 702C then generates a packet 701B and transmits the packet 701B in accordance with the adjacency information 709 and information identified in the adjacency table 319C. The forwarding engine 702C applies encapsulation information (e.g., ATM, frame relay, etc.) to the packet 701A to generate the packet 701B. The data of the packet 701B may include additional encapsulations.

[0044] Figure 7C is an exemplary diagram illustrating transmission of a packet along an LSP at a network element acting as a transient router according to one embodiment of the invention. In Figure 7C, the line card 601A receives a packet 701A that has been decapsulated from a data link layer encapsulation and that includes an ingress label and a payload. The payload includes additional encapsulations and may include additional labels. The forwarding engine 702A processes the packet 701A in accordance with the LFIB 315A. The forwarding engine 702A determines a forwarding structure and adjacency information that corresponds to the ingress label of the packet 701A in the LFIB 315A. The forwarding engine 702A may apply forwarding features (e.g., quality of service, packet classification, packet counters, etc.), which are associated with the indicated forwarding structure, to the

packet 701A. The forwarding engine 702A passes the packet 701A and the adjacency information 709 to the switching medium 703.

[0045] The switching medium 703 forwards the packet 701A and the adjacency information 709 to the line card 601C, again assuming the adjacency information 709 identifies the slot hosting the line card 601C. The forwarding engine 702C processes the packet 701A in accordance with the LSP adjacency table 316C. The forwarding engine 702C determines encapsulation information and possibly a port corresponding to the adjacency information 709. The forwarding engine 702C may apply forwarding features associated the forwarding structure indicated in the LSP adjacency table 316C. The forwarding engine 702C transmits a packet 701B that includes an egress label indicated in the LSP adjacency table 316C after encapsulating the packet 701B with data link layer encapsulation(s).

[0046] As illustrated in Figures 7A – 7C, representing LSPs with forwarding structures enables transparent application of forwarding features to packets traversing an LSP and uniform presentation of LSPs with non-LSP traffic. Forwarding features are applied to LSP traffic and non-LSP traffic with forwarding structures. The complexity of the hosting network element is reduced because forwarding features can be indexed in the same manner for LSP traffic and non-LSP traffic.

[0047] Figure 8 is a diagram of an exemplary network element according to one embodiment of the invention. In Figure 8, the network element includes a control card 803 in the control plane 301. The control card 803 is coupled with a transmission medium 805 (e.g., a system bus) in the data plane 321. The transmission medium 805 is coupled with the line cards 601A-601D. The transmission medium 805 carries information from the control card 803 to the line cards 601A-601D. The line cards 601A-601D are coupled with each other via the switching medium 803. As previously described, the switching medium may be a

separate switching unit including hardware and/or software to determine which line card to forward traffic. Alternatively, the switching medium may be a mesh.

[0048] The control card 803 and the line cards 601A – 601D illustrated in Figure 8 include memories, processors, and/or ASICs. Such memories include a machine-readable medium on which is stored a set of instructions (i.e., software) embodying any one, or all, of the methodologies described herein. Software can reside, completely or at least partially, within this memory and/or within the processor and/or ASICs. For the purpose of this specification, the term "machine-readable medium" shall be taken to include any mechanism that provides (i.e., stores and/or transmits) information in a form readable by a machine (e.g., a computer). For example, a machine-readable medium includes read only memory ("ROM"), random access memory ("RAM"), magnetic disk storage media, optical storage media, flash memory devices, electrical, optical, acoustical, or other form of propagated signals (e.g., carrier waves, infrared signals, digital signals, etc.), etc.

[0049] While the invention has been described in terms of several embodiments, those skilled in the art will recognize that the invention is not limited to the embodiments described. The method and apparatus of the invention can be practiced with modification and alteration within the spirit and scope of the appended claims. The description is thus to be regarded as illustrative instead of limiting on the invention.